Compressing and Indexing Stock Market Data

Gianni Decaroli[1] and Giovanni Manzini[1,2]

[1]Computer Science Institute, University of Eastern Piedmont, Italy
[2]IIT-CNR, Pisa, Italy

arXiv:1606.05724v2 [cs.DS] 13 Sep 2016

Abstract

We show how to build a compressed index for the offline *Order Preserving Pattern Matching* problem. Our solution is based on the new approach of decomposing the sequence to be indexed into an *order component*, containing ordering information, and a δ *component*, containing information on the absolute values. The Order Preserving Matching problem is then transformed into an exact matching problem on the order component followed by a verification phase on the δ component. Experiments show that this approach is viable and it is the first one offering simultaneously *small space usage* and *fast retrieval*.

1 Introduction

The problem of Order Preserving Pattern Matching consists in finding, inside a numerical sequence T, all subsequences whose elements are in a given relative order. For example, if the pattern is $P = (1, 2, 3, 4, 5)$ we need to find all increasing subsequences of length five; so if $T = (10, 20, 25, 30, 31, 50, 47, 49)$ we have a first match starting with the value 10, a second match starting with the value 20, and no others.

This problem is a natural generalization of the classic exact matching problem where we search for subsequences whose values are exactly those of the pattern. Order preserving matching is useful to search for trends in time series like stock market data, biomedical sensor data, meteorological data, *etc.*.[1] In the last few years this problem has received much attention. Not surprisingly, most of the results are generalization of algorithms and techniques used for exact matching. In [1, 6, 16, 18] the authors propose solutions inspired by the classical linear time Knuth-Morris-Pratt and Boyer-Moore algorithms [15, Chap. 2]. In [7] the authors consider the offline problem in which T can be preprocessed and propose an index that generalizes the classical Suffix Tree data structure [15, Chap. 5]. Finally in [2, 3, 4, 5, 9, 11] the authors consider approaches based on the concept of filtration and seminumerical matching [15, Chap. 4].

In this paper we extend to Order Preserving Matching another well known idea of exact matching: simultaneously compressing and indexing a sequence of values [20]. We show how to compactly represent a sequence T so that given a pattern P we can efficiently report all subsequences of T whose elements are in the same relative order as the elements of P. Our contribution is based on the new idea of decomposing the sequence T into two components: the *order component* and the δ *component*. Informally, the order component

[1]The title of the paper is an attempt to find an easy-to-remember name for this problem: we do not claim our results are particularly suitable to stock market data.

stores the information about the relative order of the elements of T inside a window of a preassigned size, while the δ component contains the information required for reconstructing T given the order component. The order component is stored into a compressed suffix array while the δ component is stored using an ad-hoc compression technique.

To search for a pattern we compute its ordering information and then we search it in the compressed suffix array of the order component. Since the information in the order component is only partial, this search gives us a list of potential candidates which are later verified using the δ component. In other words, the search in the compressed suffix array is a sort of filtering phase that uses the index to quickly select a set of candidates, discarding all other subsequences in T that certainly do not match. This approach can be seen as a generalization of the offline strategy in [4] as we will comment in the next section.

The overall efficiency of our approach depends on some parameters of the algorithm whose influence will be experimentally analyzed in Section 6. The bottom line is that our index takes roughly the same space as the compressor gzip and can report the order preserving occurrences of a pattern order of magnitude faster than a scan based algorithm.

2 Problem formulation and previous results

Let $T[1, n]$ denote a sequence of $n = |T|$ numerical values. We write $T[i]$ to denote the i-th element and $T[j, k]$ to denote the subsequence $T[j]T[j + 1] \cdots T[k]$. Given two sequences P, Q, we say that they are *order isomorphic*, and write $P \approx Q$, if $|P| = |Q|$ and the relative order of P's and Q's elements is the same, that is

$$P[i] \leq P[j] \iff Q[i] \leq Q[j] \qquad \text{for } 1 \leq i, j \leq |P|. \tag{1}$$

Hence $(1, 3, 4, 2)$ is order isomorphic to $(100, 200, 999, 101)$ but not to $(1, 3, 4, 5)$. Given a reference sequence $T[1, n]$ and a pattern $P[1, p]$ the *order preserving pattern matching problem* consists in finding all subsequences $T[i+1, i+p]$ such that $T[i+1, i+p] \approx P[1, p]$. In this paper we consider the offline version of the problem in which the sequence T is given in advance and we are allowed to preprocess it in order to speedup subsequent searches.

For the offline problem, the authors of [7] describe a generalization of the Suffix Tree data structure that reports the occurrences of a pattern P in $\mathcal{O}(|P| + \mathsf{occ})$ time, where occ is the number of (order-preserving) occurrences. Like the classical Suffix Tree, the one described in [7] has the drawback of using $\mathcal{O}(n \log n)$ bits of space. The only other offline solution we are aware of is the one in [4] where the authors build an index which is used as a filter to get a list of potential matches. The index used in [4] is a FM-index [10] built on the binary sequence $t_1 \cdots t_{n-1}$ such that $t_i = 1$ iff $T[i + 1] > T[i]$. The pattern to be searched is also transformed with the same rule and the corresponding binary string is searched in the index. This provides a list of potential matches that are later verified using the actual sequence T.

In this paper, following a well established research line [20], we consider the problem of designing a *compressed index* for the Order Preserving Matching problem. We describe a data structure that uses roughly 30% of the space used by the sequence T, and supports the efficient retrieval of order isomorphic subsequences. Our starting point is a generalization of the index in [4]. However, instead of extracting a binary sequence from T we use a sliding window of size q and we extract an *order component* defined over the alphabet $\{0, 1, \ldots, q - 1\}$. In addition, we compute a δ *component* containing the information not

stored in the order component. We show that these components can be compactly stored and at the same time used as an index for Order Preserving Matching.

Our solution makes use of *compressed suffix array* implemented using a FM-index [10]. The reader should be familiar with the basic concepts of this data structure, namely the fact that it works in terms of ranges of Suffix Array rows (even if the Suffix Array itself is not available) and its two basic operations: *backward search* and *LF mapping* [20].

3 Basic results

In the following we assume that all sequences contain distinct values. If this is not the case, we force this property assuming that if $i < j$ and entries $T[i]$ and $T[j]$ contain the same value then $T[i]$ is considered smaller than $T[j]$. This is the same assumption used in [16]. In some applications equal values are important and it is natural to require that equal values in the pattern correspond to equal values in T. Since our algorithms are slightly more complex for this case we postpone their discussion to Section 5.1.

Given a positive parameter q and a sequence $T[1, n]$ we define a new sequence $T_o[1, n]$ with elements over the set $\{0, 1, \ldots, q-1\}$. Let $i_q = \max(1, i-q+1)$; for $i = 1, \ldots, n$ we define

$$T_o[i] = \begin{cases} 0 & \text{if } T[i] = \min T[i_q, i] \\ k & \text{if } T[i-k] = \max\{T[j] \mid i_q \le j < i \ \wedge \ T[j] < T[i]\}. \end{cases} \tag{2}$$

In other words, we set $T_o[i] = k > 0$ when $T[i-k]$ is the immediate predecessor of $T[i]$ in the subsequence $T[i_q, i]$, that is, $T[i-k]$ is the largest element smaller than $T[i]$ in $T[i_q, i]$. If $T[i]$ is the smallest element in the sequence, then it has no predecessor and $T_o[i] = 0$.

For example, if $q = 4$ for $T = (3, 8, 3, 5, -2, 9, 6, 6)$ it is $T_o = (0, 1, 2, 1, 0, 2, 3, 1)$ (recall we are assuming $T[1] < T[3]$ and $T[7] < T[8]$). We call T_o the *order component* for T since it encodes ordering information for T's elements within a size-q window. Formally the sequence T_o depends also on q, but since we use the same q for all sequences for simplicity we omit it from the notation. Representations similar to T_o, but without a sliding window, have been used in [7, 16].

Obviously, if P and Q have the same length and $P \approx Q$ then $P_o = Q_o$. However, we are interested in finding the order preserving occurrences of P within a (much) longer reference sequence T, so we will make use of the following more general result.

Lemma 1. *Let $P[1, p]$ be any pattern and T a reference sequence. Let i be such that $P[1, p] \approx T[i+1, i+p]$. It is $T_o[i+j] = P_o[j]$ for every j such that*

$$2 \le j \le p \qquad \textit{and} \qquad (j - T_o[i+j]) \ge 1.$$

Proof. Let j be such that the above condition is satisfied and let $w = \min(q, i+j)$ and $v = \min(q, j)$. Note that w (resp. v) is the size of the subsequence which is considered for determining $T_o[i+j]$ (resp. $P_o[i+j]$). Clearly $w \ge v$.

If $T_o[i+j] = 0$, then $T[i+j]$ is the smallest element in the subsequence $T[i+j-w+1, i+j]$. A fortiori $T[i+j]$ is the smallest element in $T[i+j-v+1, i+j]$. The hypothesis $P[1, p] \approx T[i+1, i+p]$ implies that likewise $P[j]$ must be the smallest element in $P[j-v+1, j]$ so it must be $P_o[j] = 0$.

Assume now $T_o[i+j] > 0$ and let $\ell_j = j - T_o[i+j]$. By construction $T[i+\ell_j]$ is the immediate predecessor of $T[i+j]$ in the subsequence $T[i+j-w+1, i+j]$. The condition

$j - T_o[i+j] \geq 1$ implies $1 \leq \ell_j < j$. We want to show that $P_o[j] = T_o[i+j]$, that is, $P[\ell_j]$ is the immediate predecessor of $P[j]$ in $P[j-v+1, j]$.

Since $P[1,p] \approx T[i+1, i+p]$ and $T[i+\ell_j] < T[i+j]$ we have $P[\ell_j] < P[j]$. To complete the proof we need to show that there is no k, $j-v < k < j$, such that $P[\ell_j] < P[k] < P[j]$. But if this were the case we would have $T[i+\ell_j] < T[i+k] < T[i+j]$ contrary to the assumption that $T[i+\ell_j]$ is the immediate predecessor of $T[i+j]$ in $T[i+j-w+1, i+j]$. □

Since it is always $T_o[i+j] \leq q-1$, the above lemma tells us that if $P[1,p] \approx T[i+1, i+p]$ we must have $T_o[i+j] = P_o[j]$ for every j, $q \leq j \leq p$. For $j = 2, 3, \ldots, q-1$ the lemma establishes that $T_o[i+j]$ must be either equal to $P_o[j]$ or larger than $j-1$. Summing up we have the following corollary.

Corollary 2. *Given a reference $T[1,n]$ and a pattern $P[1,p]$ with $p > q$, if $P[1,p] \approx T[i+1, i+p]$ then we must have*

$$T_o[i+2, i+p] = x_2\, x_3\, \cdots x_{q-1}\, P_o[q]\, P_o[q+1] \cdots P_o[p] \tag{3}$$

where for $j = 2, \ldots, q-1$ either $x_j = P_o[j]$ or $x_j \geq j$. □

In view of the above corollary, our strategy to solve the order preserving matching problem is to build a compressed index for the sequence T_o. Then, given a pattern $P[1,p]$ with $p > q$, we compute $P_o[1,p]$ and then the set of positions $i_1, i_2, \ldots, i_m$ satisfying (3). Clearly, we can have $P[1,p] \approx T[i+1, i+p]$ only if $i \in \{i_1, i_2, \ldots, i_m\}$. However, since Corollary 2 states only a necessary condition, some of these positions could be false positives and a verification step is necessary. For the verification step we need the actual values of the sequence T. Since we are interested in indexing *and* compressing T, instead of the original representation we save space using a representation that takes advantage of the values in T_o that are stored in the index.

Given $T[1,n]$ and the corresponding ordering component $T_o[1,n]$, we define a new sequence $T_\delta[1,n]$ as follows. Let $T_\delta[1] = T[1]$. For $i = 2, \ldots, n$ let $i_q = \max(1, i-q+1)$ and:

$$T_\delta[i] = \begin{cases} \min T[i_q, i-1] - T[i] & \text{if } T_o[i] = 0 \\ T[i] - T[i - T_o[i]] & \text{if } T_o[i] > 0. \end{cases} \tag{4}$$

Observe that for $i \geq 2$, $T_\delta[i] \geq 0$. Indeed, if $T_o[i] = 0$ then by (2) $T[i] \leq \min T[i_q, i-1]$. If $T_o[i] > 0$ since $T[i - T_o[i]]$ is the immediate predecessor of $T[i]$ in $T[i_q, i]$ it is $T_\delta[i] \geq 0$.

The sequence T_δ is called the δ *component* of T. While T_o provides information on the ordering of T's elements, T_δ contains information on the absolute values of T's elements. Together these two sequences contain the same information as T, as shown by the following lemma.

Lemma 3. *Given $T_o[1,n]$ and $T_\delta[1,n]$ it is possible to retrieve $T[1,n]$ in linear time.*

Proof. It is $T[1] = T_\delta[1]$. For $i = 2, \ldots, n$ let $i_q = \max(1, i-q+1)$. From (4) we get

$$T[i] = \begin{cases} \min T[i_q, i-1] - T_\delta[i] & \text{if } T_o[i] = 0, \\ T_\delta[i] + T[i - T_o[i]] & \text{if } T_o[i] > 0. \end{cases}$$

□

Summing up, our approach to compress and index an integer sequence T and support order preserving pattern matching is the following:

1. Select a window size q and build the ordering component T_o and the δ component T_δ.
2. Build a compressed full-text index for T_o and compress T_δ taking advantage of T_o.

The compressed index for T_o and the compressed representation of T_δ constitute our index. To search the order preserving occurrences of a pattern of length at least q using our index:

1. Given $P[1, p]$ compute the corresponding ordering component P_o.
2. Use the full-text index for T_o to find the set of candidate positions $i_1, i_2, \ldots, i_m$ satisfying Corollary 2.
3. For each candidate position i_k use T_o and T_δ to retrieve $T[i_k + 1, i_k + p]$ and verify whether $P[1, p] \approx T[i_k + 1, i_k + p]$.

The above description is quite general and can be realized in many different ways. In the following sections we describe our particular implementation and experimentally measure its effectiveness.

4 Representation of the components T_o and T_δ

We represent $T_o[1, n]$ using a Compressed Suffix Array (csa) consisting of a Huffman shaped Wavelet Tree [14] built on the BWT of the sequence T_o (in our experiments we used the csa_wt class from the sdsl-lite library [12]). This approach guarantees a reasonable compression of T_o and supports very efficiently the search inside T_o. More precisely, given a pattern p the above csa computes in $\mathcal{O}(p)$ time the range of rows $[b, e]$ of the Suffix Array of T_o which are prefixed by p (see [10, 20] for details). To find the actual position in T_o of each occurrence of p, the csa stores the set of Suffix Array entries containing text positions which are multiple of a previously chosen block size B. Then, for each row $r \in [b, e]$ we walk backward in the text using the LF-map until we reach a marked Suffix Array entry from which we can derive the position in T_o of the occurrence that prefixes row r. The above scheme uses $\mathcal{O}(n + (n/B) \log n)$ bits of space and can find the position of all (exact) occurrence of p in T_o in $\mathcal{O}(|p| + B\,\mathsf{occ})$ time, where $\mathsf{occ} = e - b + 1$ is the number of occurrences. Clearly, the parameter B offers a trade-off between space usage and running time.

Having chosen a representation for T_o we now consider the problem of compactly storing the information in $T_\delta[1, n]$. We do not need to search inside T_δ however, during the verification phase, we do need to extract (decompress) the values in random portions of T_δ. For this reason we split T_δ in blocks of size B (ie the same size used for the blocks in the csa of T_o) and we compress each block independently. The k-th block consists of the subsequence $T_\delta[kB + 1, kB + B]$, except for the last block which has size $n \bmod B$. Additionally, we use a header storing the starting position of each block. Hence, given a block number we can decompress it in $\mathcal{O}(B)$ time.

To compactly represent a block of T_δ we take advantage of the fact that the corresponding values in T_o are available during compression and decompression. Recalling the definition of $T_\delta[i]$ in (4), we partition the values in T into three classes:

1. those such that $T[i] = \min T[i_q, i]$ are called *minimal*;
2. those such that $T[i] = \max T[i_q, i]$ are called *maximal*;

Name	# Values	Description
ibm	2,167,147	ibm stock tick data from Sep. 2011 to Feb. 2016
prices	31,559,990	daily, hourly, and 5min US stock prices
tmax	15,476,000	max daily temperature from 424 US stations
ecg	20,140,000	22 hours and 23 minutes of ECG data
rwalk	50,000,000	random walk with integer steps in the range $[-20, 20]$
rand	50,000,000	random integers in the range $[-20, 20]$

Figure 1: Files used in our experiments. All values are 32-bit integers so the size in bytes of the files is four times the number of values. All test files and the source code of our algorithms are available at `https://people.unipmn.it/manzini/stockmarket/`.

3. all other values are called *intermediate*.

The class of a value can be easily determined by both compressor and decompressor: minimal values are those such that $T_\delta[i] = 0$; maximal values are those such that $T_\delta[i] \neq 0$ and $T[i - T_\delta[i]] = \max T[i_q, i-1]$; all other values are intermediate. For each block we define

$$m = \max\{T_\delta[i] \mid i \text{ is minimal}\}, \qquad M = \max\{T_\delta[i] \mid i \text{ is maximal}\};$$

and we store these two values at the beginning of the block. When we encounter a minimal (resp. maximal) value $T[i]$ we know that the corresponding value $T_\delta[i]$ will be in the range $[0, m+1)$ (resp. $[0, M+1)$). The interval is semi-open since the right extreme is excluded. When we encounter a intermediate value $T[i]$ we know that $T_\delta[i]$ will be in the range $[0, v - T[i - T_o[i]])$ where v is the smallest element in $T[i_q, i-1]$ largest than $T[i - T_o[i]]$ (note that if $v = T[i - T_o[i]] + 1$ it is $T_\delta[i] = 0$ and there is nothing to encode).

Summing up, compressing a block of the sequence T_δ amounts to compressing a sequence of non-negative integers $\ell_1, \ell_2, \ldots, \ell_B$ with the additional information that for $i = 1, \ldots, B$, $\ell_i < w_i$ where the values $w_1, \ldots, w_B$ are known during both compression and decompression.[2] Let $\mathsf{bin}(k)$ denote the binary representation of the integer k. We have experimented with three different integer encoders.

lsk: Log-skewed coding described in [19] (but probably known for a long time). Encodes an integer $\ell \in [0, w)$ using at most $|\mathsf{bin}(w-1)|$ bits. If w is not a power of two the smallest values in the range $[0, w)$ are encoded using less than $|\mathsf{bin}(w-1)|$ bits. The encoding takes at most $\log(w) + \mathcal{O}(1)$ bits.

dlt: Delta coding [8]. Encodes an integer $\ell \in [0, w)$ using $2|\mathsf{bin}(\ell+1)| - 1$ bits consisting of the unary representation of $|\mathsf{bin}(\ell+1)|$ followed by the bits of $\mathsf{bin}(\ell+1)$ except for the most significant bit. This encoding takes $2\log(\ell+1) + \mathcal{O}(1)$ bits, is efficient for small ℓ but does not take advantage of w.

lsd: Log-skewed-delta coding: a new approach combining lsk and dlt. Encodes $\ell \in [0, w)$ by first using lsk to encode the integer $|\mathsf{bin}(\ell+1)| \in [0, |\mathsf{bin}(w)|]$ followed by the bits of $\mathsf{bin}(\ell+1)$ except for the most significant bit. The encoding takes at most $\log(\ell+1) + \log\log w + \mathcal{O}(1)$ bits.

We have tested the above three coders for different values of the window size q and of the block size B on the collection of test files described in Fig. 1. The results for $B = 32$ are shown in Table 1 together with the space occupancy of the csa for the sequence T_o. We can summarize this first set of experiments as follows:

1. Among the three encoders lsk is the one that compresses better, closely followed by lsd, while dlt is clearly worse. We have also tested a mixed approach in which intermediate values are encoded with lsk while minimal and maximal values are encoded with lsd, but this strategy did not improve over the simpler lsk encoder. For this reason in rest of the paper delta values are always encoded with lsk.

2. As the window size q increases the cost of storing T_δ decreases while the csa size increases. This was to be expect since a larger q means that more information is contained in T_o.

Summing up, for a given block size B and window size q our "stock market index" (smi from now on, again with no claim that it is particularly suitable for stock market data) consists of the csa for T_o and a compressed representation of T_δ obtained using the lsk encoder. Table 2 shows the overall space of our index for different values of B and q compared to the space used by gzip and by the state of the art compressor xz. We can see that smi's space usage is essentially at par with gzip's: it can be smaller or larger depending on the block size B. As expected xz compression is clearly superior to both. These data show that smi uses an amount of space similar to modern compressed full text indices for exact pattern matching [13].

5 Searching a Stock Market Index

Given a pattern $P[1, p]$, to find its order-preserving occurrences in T we compute its order component P_o and use Corollary 2. For example, if $q = 4$ and $P_o = (0, 0, 1, 3, 2, 0)$, Corollary 2 tells us that in correspondence to $P_o[4, p] = (3, 2, 0)$ T_o must be $(3, 2, 0)$ (an exact match), while in correspondence of $P_o[3] = 1$, T_o must be either 1 or 3, and in correspondence of $P_o[2] = 0$, T_o must be either 0, 2, or 3. In other words if $T[i+1, i+p] \approx P$ then $T_o[i + 2, i + p]$ must be one of the following 6 sequences:

$$(0, 1, 3, 2, 0) \quad (0, 3, 3, 2, 0) \quad (2, 1, 3, 2, 0) \quad (2, 3, 3, 2, 0) \quad (3, 1, 3, 2, 0) \quad (3, 3, 3, 2, 0).$$

In general, if the window size is q the number of sequences satisfying Corollary 2 is $(q-1)!$. Clearly, not all the $(q - 1)!$ subsequences will necessarily occur inside T_o. Using T_o's csa we compute the range of Suffix Array rows which are prefixed by each one of the $(q - 1)!$ subsequences mentioned above. Recall that the basic operation of a csa is the *backward search* in which, given the range of rows prefixed by a substring α and a character c, we find in $\mathcal{O}(1)$ time the range of rows prefixed by $c\alpha$. This suggests to compute the desired set of row ranges with a two steps procedure: first (Phase 1) with $p-q+1$ backward search steps we compute the range of rows prefixed by $P[q, p]$; then (Phase 2) with additional backward search steps we compute the range of rows prefixed by $x_2x_3 \cdots x_{q-1}P[q] \cdots P[p]$ for each $(q - 2)$-tuple $x_2, \ldots, x_{q-1}$ satisfying the conditions of Corollary 2. Phase 2 can

[2]While T_δ blocks are compressed independently, T_o values are defined globally. Hence for the first $q-1$ values of a block T_o may point to a value in the previous block. In this case we store in T_δ also the position of the predecessor inside the block.

		ibm	prices	tmax	ecg	rwalk	rand
$B=32$ $q=4$	csa	0.09	0.10	0.10	0.09	0.10	0.11
	lsk	0.18	0.25	0.22	0.18	0.24	0.24
	dlt	0.19	0.30	0.25	0.20	0.28	0.28
	lsd	0.18	0.26	0.23	0.19	0.26	0.26
$B=32$ $q=8$	csa	0.11	0.13	0.13	0.11	0.13	0.14
	lsk	0.17	0.24	0.20	0.17	0.23	0.22
	dlt	0.18	0.29	0.23	0.20	0.27	0.25
	lsd	0.17	0.25	0.21	0.18	0.25	0.23
$B=32$ $q=12$	csa	0.12	0.14	0.15	0.13	0.15	0.16
	lsk	0.17	0.23	0.19	0.17	0.22	0.21
	dlt	0.18	0.28	0.23	0.19	0.27	0.25
	lsd	0.17	0.25	0.21	0.18	0.24	0.22

Table 1: Space usage of T_o's csa and of T_δ compressed with three different encoders: log-skewed (lsk), delta coding (dlt), and log-skewed delta (lsd). The reported values are the ratio between the size of compressed file over the size of the test file (both expressed in bytes).

		ibm	prices	tmax	ecg	rwalk	rand
$B=32$	smi $q=4$	0.26	0.35	0.32	0.27	0.34	0.35
	smi $q=8$	0.27	0.37	0.33	0.29	0.36	0.36
	smi $q=12$	0.28	0.38	0.34	0.30	0.37	0.37
$B=64$	smi $q=4$	0.20	0.29	0.26	0.21	0.28	0.28
	smi $q=8$	0.21	0.30	0.26	0.22	0.29	0.28
	smi $q=12$	0.21	0.31	0.27	0.23	0.30	0.29
$B=96$	smi $q=4$	0.18	0.27	0.24	0.19	0.25	0.26
	smi $q=8$	0.19	0.28	0.24	0.20	0.27	0.26
	smi $q=12$	0.19	0.28	0.25	0.21	0.28	0.26
	gzip --best	0.22	0.37	0.24	0.19	0.36	0.24
	xz --best	0.13	0.24	0.17	0.12	0.23	0.18

Table 2: Space usage of the Stock Market Index (smi) for different values of B and q. Each value is the ratio between the size of the index over the size of the test file, both expressed in bytes. For completeness we show also the space usage for gzip and xz.

require up to $q!$ backward search steps, but the number of steps is also upper bounded by q times the number of row ranges obtained at the end of the phase, which is usually much smaller.

At the end of Phase 2 we are left with a set of rows each one representing a position in T where an order preserving match can occur. To verify if there is actually a match we have to decompress the corresponding subsequence of T and compare it with P. This is done (Phase 3) taking again advantage of the properties of T_o's csa. Given a row index r representing a position in T_o prefixed by a string $x_2x_3\cdots x_{q-1}P[q]\cdots P[p]$ we use the LF-map to move backward in T_o until we reach a marked position, that is, a position in T_o (and hence in T) which is a multiple of the block size B (say position ℓB) and marks the beginning of block ℓ. Each time we apply the LF-map we also obtain a symbol y_i of T_o hence when we reach the beginning of the block we also have the sequence

$$y_1\, y_2\, \cdots y_k\, x_2\, x_3\, \cdots x_{q-1}P[q]\cdots P[p]$$

of T_o values from the beginning of the block till the position corresponding to $P[p]$. Using this information and the compressed representation of T_δ (whose blocks are compressed independently) we are able to retrieve the corresponding T values

$$T[\ell B+1]T[\ell B+2]\cdots T[\ell B+k]T[\ell B+k+1]\cdots T[\ell B+k+p-1]$$

and determine if there is an actual order preserving match between the last p values of the above sequence and $P[1,p]$. If there is match the algorithm outputs the starting position in T of the matching sequence (the value $\ell B+k$ in the above example).

Phase 3 is usually the most expensive step since the algorithm has to consider one candidate at a time and for each candidate we need to reach the beginning of the block containing it. We can therefore expect that its running time is linearly affected by the block size B. Note that in our implementation Phase 2 and 3 are interleaved: as soon as we have determined a range of rows prefixed by one of the patterns in Corollary 2 we execute Phase 3 for all rows in the range before considering any other row range.

5.1 Strict index: Taking care of equal values

In this section we discuss the modifications which must be made to our index and algorithms in the general case when we cannot assume all values in our sequences are distinct. If equal values can occur, according to definition (1), to have $P[1,p]\approx T[i+1,i+p]$ whenever $P[j]=P[k]$ we must have $T[i+j]=T[i+k]$. In Section 3 we side-stepped this problem assuming that later occurrences of a value are greater that previous occurrences: Under that assumption, if $P[j]=P[k]$ with $j<k$ it is acceptable also that $T[i+j]<T[i+k]$.

Since, if $P[1,p]\approx T[i+1,i+p]$ according to the stricter definition they are order isomorphic also under the assumption used in Section 3, we can build an index for T as described above and only change the verification of the single candidates (Phase 3). However, there is a non-trivial improvements that can be made to the algorithm to reduce the number of candidates. The improvement is a consequence of the following lemma.

Lemma 4. *In the presence of equal values, if $P[1,p]\approx T[i+1,i+p]$ and $P[j-P_o[j]]=P[j]$ with $P_o[j]>0$, then it is $T_o[i+j]=P_o[j]$.*

Proof. If $P_o[j]>0$ and $P[j-P_o[j]]=P[j]$ then $P[j-P_o[j]]$ must be rightmost value to the left of $P[j]$ which is equal to $P[j]$. The hypothesis $P[1,p]\approx T[i+1,i+p]$ implies that the same should be true for $T[i+j-P_o[j]]$ with respect to $T[i+j]$. Hence we must have $T_o[i+j]=P_o[j]$ as claimed. □

		ibm	prices	tmax	ecg	rwalk	rand
$B = 32$	ssmi $q=4$	0.27	0.35	0.32	0.28	0.34	0.35
	ssmi $q=8$	0.28	0.37	0.34	0.30	0.36	0.36
	ssmi $q=12$	0.29	0.39	0.35	0.31	0.38	0.37
$B = 64$	ssmi $q=4$	0.20	0.29	0.26	0.21	0.27	0.28
	ssmi $q=8$	0.21	0.31	0.27	0.23	0.29	0.29
	ssmi $q=12$	0.22	0.32	0.28	0.24	0.30	0.30
$B = 96$	ssmi $q=4$	0.18	0.27	0.24	0.19	0.25	0.26
	ssmi $q=8$	0.19	0.29	0.25	0.21	0.27	0.26
	ssmi $q=12$	0.20	0.30	0.26	0.22	0.28	0.27
	gzip --best	0.22	0.37	0.24	0.19	0.36	0.24
	xz --best	0.13	0.24	0.17	0.12	0.23	0.18

Table 3: Space usage of the Strict Stock Market Index (ssmi) for different values of B and q. Each value is the ratio between the size of the index over the size of the test file, both expressed in bytes. For completeness we show also the space usage for gzip and xz.

Because of the above Lemma, in the presence of equal values Corollary 2 should be modified so that for $j = 2, \ldots, q-1$ we should consider the case $x_j > j$ only if $P_o[j] = 0$ or $P[j - P_o[j]] \neq P[j]$. Indeed, by Lemma 4 whenever $P_o[j] \neq 0$ and $P[j - P_o[j]] = P[j]$ we must have $T_o[i + j] = P_o[j]$. Hence, in the presence of equal values in the pattern we need to do consider a smaller number of candidates (see the discussion of Phase 2 of the search algorithm at the beginning of Section 5).

If we are only interested to search in strict mode, ie considering equal values, we can take advantage of the presence of equal values in T to completely avoid the encoding of some of the δ components. This will yield a smaller compressed file and faster search. To achieve this result we slightly change the definition of the order component given in (2) as follows. When computing $T_o[i]$, assume the immediate predecessor of $T[i]$ in $T[i_q, i-1]$ is some value y that appears at least twice in $T[i_q, i-1]$. Say $y = T[k_1] = T[k_2] = \cdots = T[k_h]$ with $i_q \leq k_h < \cdots k_2 < k_1 \leq i-1$. If this is the case, if $T[i] = y$ we set as usual $T_o[i] = i-k_1$; whereas if $T[i] > y$ we set $T_o[i] = i - k_2$ (ie we skip an occurrence of y).

Using this scheme, whenever $T_o[i] = k > 0$ and there is another occurrence of $T[i-k]$ in $T[i_q, i-1]$, both the coder and the decoder can tell whether $T[i] = T[i-k]$ or $T[i] > T[i-k]$ depending on the position of the other occurrence of $T[i-k]$. In particular, if both the coder and decoder know that $T[i] = T[i-k]$ there is no need to compute and store the corresponding value $T_\delta[i]$. The space saving for T_δ can be significant when T contain long stretches of equal values. If the modified procedure is used to compute both T_o and P_o the search algorithm does not need to be modified.

A completely different approach to search in strict mode is to add an additional bit to each nonzero entry of T_o to distinguish whether $T[i - T_o[i]]$ is equal to $T[i]$ or not. More precisely, we define the order component as follows:

$$T_o[i] = \begin{cases} 0 & \text{if } T[i] < \min T[i_q, i-1] \\ 2k & \text{if } T[i-k] = \max\{T[j] \mid i_q \leq j < i \wedge T[j] \leq T[i]\},\ T[i-k] \neq T[i]. \\ 2k-1 & \text{if } T[i-k] = \max\{T[j] \mid i_q \leq j < i \wedge T[j] \leq T[i]\},\ T[i-k] = T[i]. \end{cases}$$

With this approach T_o contains values in the range $[0, 2q-1]$ so we expect its index to take

more space. However, we also expect T_δ to takes less space since values corresponding to odd entries in T_o don't have to be encoded. We have implemented this variant, from now on called the *strict index*, and its space usage is shown in Table 3. Comparing Table 3 with Table 2 we see that for the same values B and q the strict index takes only slightly more space, the difference being about 1% of the original (uncompressed) file size.

6 Experimental results

All tests have been performed on a desktop PC with eight Intel-I7 3.40GHz CPUs running Linux-Debian 8.3. All tests used a single CPU. Note that Phase 3 of our algorithm can be easily parallelized using multiple threads (and to some extent also Phase 2), but we leave this development to future research. All tests involved 1000 patterns of length 15, 20, and 50 extracted from the same file where the patterns are later searched, so every search reports at least one occurrence. The patterns where extracted selecting 1000 random position in the in file. Note that patterns occurring more often are more likely to be selected so this setting is the least favorable for our algorithm: like all algorithms based on an index, it is much faster when the pattern does not occur, or occurs relatively few times.

6.1 Filtering effectiveness

Since Phases 1 and 2 of our algorithms produce a set of candidates that must be verified in Phase 3, in our first experiment we measure how effective are the first two Phases in producing only a small number of candidates which are later discarded (that is, how effective are Phases 1 and 2 in producing a small number of false positives). The results of this experiment for smi are reported in Table 4. We see that for patterns of length 15 the average number of occurrences is surprisingly high, with a peak of 73000+ average occurrences for the file ecg (recall the random selection of patterns favors those which occur more often). Despite this, we see that in most cases for Phase 2 the ratio is smaller than 1.50, that is, the number of candidates at the end of Phase 2 is less than one and a half times the number or actual occurrences. Note that the poor performance of Phase 1 for $q = 12$ was to be expected since when $|P| = 15$ Phase 1 will only search in T_o's csa the last 4 symbols of P_o which cannot provide an effective filter.

Not surprisingly for patterns of length 20 and 50 both phases are much more effective. In particular, the number of false positives at the end of Phase 2 is usually smaller than 10% the number of occurrences (ratio below 1.10). Note that in terms of filtering power there is not a clear winner among the different window sizes. The reason is that a larger window implies that the information stored in T_o is "more accurate" but at the same time Phase 2 will start earlier and be less effective since it will generate $(q-1)!$ row ranges.

For the search in strict mode the number of false positives as a function of the window size q follows a similar trend. The data are shown in Table 7 in the Appendix.

6.2 Running times

Our search algorithms have no direct competitors since they are the first ones to combine compression and indexing. We expect them to be faster than scan based (online) algorithms [2, 6, 16, 18, 5], at least for sufficiently long files, and slower than the offline algorithms [4, 7] designed without considering the problem of compressing the reference

$\|P\| = 15$		ibm	prices	tmax	ecg	rwalk	rand
ave # occs		1155.43	293.83	1469.18	73054.84	4.36	1.00
$q = 4$	Phase 1	2.22	1.97	1.17	1.33	15.87	13.07
	Phase 2	1.23	1.21	1.01	1.06	6.45	6.24
$q = 8$	Phase 1	7.00	8.99	1.58	2.46	139.62	5.79
	Phase 2	1.21	1.19	1.01	1.05	3.52	1.09
$q = 12$	Phase 1	42.73	170.71	5.20	7.75	13745.64	2620.21
	Phase 2	1.28	1.31	1.02	1.06	13.32	1.79

$\|P\| = 20$		ibm	prices	tmax	ecg	rwalk	rand
ave # occs		273.43	55.45	1168.09	24299.10	1.01	1.00
$q = 4$	Phase 1	1.69	1.17	1.05	1.21	1.40	1.02
	Phase 2	1.11	1.05	1.00	1.05	1.19	1.01
$q = 8$	Phase 1	4.95	3.42	1.41	2.31	1.92	1.00
	Phase 2	1.14	1.08	1.01	1.04	1.02	1.00
$q = 12$	Phase 1	14.38	9.66	1.56	4.57	89.67	1.01
	Phase 2	1.20	1.11	1.01	1.06	1.07	1.00

$\|P\| = 50$		ibm	prices	tmax	ecg	rwalk	rand
ave # occs		1.00	1.00	380.55	5.45	1.00	1.00
$q = 4$	Phase 1	2.14	1.00	1.02	2.05	1.00	1.00
	Phase 2	1.08	1.00	1.00	1.78	1.00	1.00
$q = 8$	Phase 1	2.47	1.00	1.07	1.77	1.00	1.00
	Phase 2	1.00	1.00	1.00	1.14	1.00	1.00
$q = 12$	Phase 1	4.86	1.00	1.12	2.53	1.00	1.00
	Phase 2	1.00	1.00	1.01	1.05	1.00	1.00

Table 4: False positives as a function of the window size q. The first row shows the average number of actual occurrences for the patterns in the test set. The other rows show the ratios between candidates and actual occurrences at the end of Phase 1 and 2.

sequence. We plan a comprensive experimental analysis in a future study; in this section we are mainly interested in optimizing our algorithms and understanding the influence of the window size q and of the block size B on the running times.

Table 5 reports the running times of our stock market index (smi) for different values of q and B. As a reference we report also the running times of a simple scan based search algorithm in which we check each text position with the verification algorithm outlined in [5, Sec. 3] (this is the same verification algorithm used by Phase 3 of our algorithm). We observe that for the largest files (prices, rwalk, rand) smi is at least two order of magnitudes faster than scan, the difference being more evident when the average number of occurrences is small. On the other hand, for the file ecg with $|P| = 15$ there are 73000+ average occurrences per pattern, and our algorithm is slower than scan. The explanation is that just to extract 73000 length-15 subsequences smi decompresses more than one million input values which is roughly 1/20 of the entire file. Since subsequences are extracted in random order and there are additional overheads due to compression, it is

$\|P\| = 15$		ibm	prices	tmax	ecg	rwalk	rand
ave # occs		1155.43	293.83	1469.18	73054.84	4.36	1.00
$B=32$	$q=4$	3.64	0.94	3.01	220.34	0.15	0.04
	$q=8$	5.09	1.38	4.50	318.10	0.15	0.02
	$q=12$	7.34	2.32	5.77	399.17	1.09	0.08
$B=64$	$q=4$	6.49	1.71	6.01	407.17	0.26	0.07
	$q=8$	9.19	2.54	8.99	604.65	0.25	0.03
	$q=12$	12.88	4.02	11.55	768.24	1.55	0.11
scan		17.28	206.44	101.90	129.71	300.89	352.94

$\|P\| = 20$		ibm	prices	tmax	ecg	rwalk	rand
ave # occs		273.43	55.45	1168.09	24299.10	1.01	1.00
$B=32$	$q=4$	0.78	0.14	2.44	77.52	0.01	0.01
	$q=8$	1.14	0.22	4.13	114.90	0.02	0.02
	$q=12$	1.64	0.29	4.73	143.78	0.03	0.03
$B=64$	$q=4$	1.39	0.26	4.69	139.45	0.02	0.02
	$q=8$	2.04	0.39	7.24	210.68	0.03	0.03
	$q=12$	2.93	0.51	9.79	268.84	0.04	0.04
scan		17.10	208.34	101.65	127.16	303.01	367.82

$\|P\| = 50$		ibm	prices	tmax	ecg	rwalk	rand
ave # occs		1.00	1.00	380.55	5.45	1.00	1.00
$B=32$	$q=4$	0.01	0.02	0.91	0.04	0.02	0.02
	$q=8$	0.02	0.03	1.44	0.05	0.03	0.03
	$q=12$	0.03	0.04	1.92	0.07	0.04	0.04
$B=64$	$q=4$	0.01	0.02	1.52	0.07	0.02	0.03
	$q=8$	0.02	0.03	2.40	0.07	0.04	0.04
	$q=12$	0.03	0.04	3.20	0.09	0.05	0.05
scan		16.76	202.29	100.38	120.99	296.46	380.01

Table 5: smi vs scan: Average running times, in milliseconds, for searching 1000 random patterns of length 15, 20, and 50. Running times do not include the time to load the compressed index (for smi) or the uncompressed text (for scan).

$\|P\| = 15$		ibm	prices	tmax	ecg	rwalk	rand
ave # occs		92.63	59.70	1392.66	2142.49	2.38	1.00
$B=32$	$q=4$	0.23	0.17	3.04	7.73	0.06	0.02
	$q=8$	0.25	0.21	3.80	9.39	0.09	0.03
	$q=12$	0.29	0.44	3.77	11.98	0.72	0.05
$B=64$	$q=4$	0.43	0.31	6.39	14.26	0.11	0.03
	$q=8$	0.49	0.41	8.01	18.89	0.13	0.03
	$q=12$	0.57	0.64	8.65	23.65	0.88	0.06
scan		16.47	252.64	122.92	135.78	406.92	419.95

$\|P\| = 20$		ibm	prices	tmax	ecg	rwalk	rand
ave # occs		24.22	1.49	988.39	106.78	1.00	1.00
$B=32$	$q=4$	0.06	0.02	2.15	0.47	0.01	0.01
	$q=8$	0.07	0.03	2.73	0.56	0.03	0.03
	$q=12$	0.08	0.04	2.70	0.81	0.05	0.04
$B=64$	$q=4$	0.11	0.02	4.26	0.85	0.02	0.02
	$q=8$	0.13	0.04	5.54	1.07	0.03	0.04
	$q=12$	0.17	0.05	5.82	1.46	0.06	0.05
scan		16.59	246.37	121.52	132.68	403.77	410.79

$\|P\| = 50$		ibm	prices	tmax	ecg	rwalk	rand
ave # occs		1.00	1.00	371.99	1.00	1.00	1.00
$B=32$	$q=4$	0.02	0.02	0.87	0.02	0.02	0.02
	$q=8$	0.02	0.03	1.01	0.03	0.04	0.04
	$q=12$	0.03	0.05	0.97	0.04	0.06	0.06
$B=64$	$q=4$	0.01	0.02	1.55	0.02	0.02	0.03
	$q=8$	0.02	0.04	1.90	0.03	0.04	0.05
	$q=12$	0.03	0.06	1.87	0.04	0.06	0.07
scan		15.92	241.06	114.22	131.76	392.10	345.17

Table 6: ssmi vs scan: Average running times, in milliseconds, for searching 1000 random patterns of length 15, 20, and 50. Running times do not include the time to load the compressed index (for ssmi) or the uncompressed text (for scan).

not surprising that a simple scan is faster. This suggests that to efficiently deal with this almost pathological number of occurrences it is necessary to use ad hoc techniques as it was done for other indices [17, 21]. The results for the two smallest files in our collection (ibm 2 million values, tmax 15 million values) show that our approach clearly outperforms scan also when the number of occurrences is relatively large and the size of the input is relatively small.

The results in Table 5 show that doubling the block size B from 32 to 64 usually doubles smi's running time which is and indirect confirmation that Phase 3 (the only phase influenced by B) is the most time consuming. We also see that the smallest running times are obtained consistently with $q = 4$, the only exception being rwalk and rand for $|P| = 15$. The likely explanation here is that for a smaller q all operations on the Wavelet Tree underlying the csa are faster and this makes up for the better filtration obtained in some instances by Phase 1 and 2 for larger values of q. Since, according to Table 1, $q = 4$ also gives the highest compression, among the tested values $q = 4$ appears to be the optimal choice. We have tested smi also for $q = 3$ (smi requires $q > 2$). We do not report the results here, but it turns out that $q = 3$ yields essentially the same compression as $q = 4$ and very similar search speed (there is not a clear winner among the two).

Table 6 shows the search performance of our strict stock market index (ssmi). Since the order preserving matching is strict, the average number of occurrences of each pattern is much smaller (up to 200 times for ecg for $|P| = 15$) and is at most 2142. Hence, we do not have the pathological cases we had previously and ssmi was always faster than scan. For $B = 32$ and $q = 4$ the minimum speedup was 17 (for ecg and $|P| = 15$). For ssmi we see that $q = 4$ always gives the smallest space usage and the fastest speed, and that doubling the block size doubles the running times, showing again that also for ssmi the cost of Phase 3 dominates the overall cost.

7 Concluding Remarks

In this paper we have proposed a compressed index for the order preserving pattern matching problem, considering both the weak and the strict order preserving matching models. Our approach is based on the new idea of splitting the original sequence into two complementary components: the order component and the δ component. The problem of finding the order preserving occurrences of a pattern is transformed into an exact search problem on the order component followed by a verification phase using the δ component. Experiments show that our index has a space usage similar to gzip and can find order preserving occurrences much faster than a sequential scan.

Our approach is quite general and improvements could be obtained by changing some implementation choices. For example, we index the order component using a Wavelet-Tree based FM-index; to improve the performances for inputs with many (order preserving) repetitions we can use a different compressed full-text index. The compression of the δ component can also be radically changed without altering the overall scheme.

References

[1] Djamal Belazzougui, Adeline Pierrot, Mathieu Raffinot, and Stéphane Vialette. Single and multiple consecutive permutation motif search. In *ISAAC*, volume 8283 of *Lecture Notes in Computer Science*, pages 66–77. Springer, 2013.

[2] Domenico Cantone, Simone Faro, and M. Oguzhan Külekci. An efficient skip-search approach to the order-preserving pattern matching problem. In *Stringology*, pages 22–35. Department of Theoretical Computer Science, Czech Technical University in Prague, 2015.

[3] Tamanna Chhabra, Emanuele Giaquinta, and Jorma Tarhio. Filtration algorithms for approximate order-preserving matching. In *SPIRE*, volume 9309 of *Lecture Notes in Computer Science*, pages 177–187. Springer, 2015.

[4] Tamanna Chhabra, M. Oguzhan Külekci, and Jorma Tarhio. Alternative algorithms for order-preserving matching. In *Stringology*, pages 36–46. Department of Theoretical Computer Science, Czech Technical University in Prague, 2015.

[5] Tamanna Chhabra and Jorma Tarhio. A filtration method for order-preserving matching. *Inf. Process. Lett.*, 116(2):71–74, 2016.

[6] Sukhyeun Cho, Joong Chae Na, Kunsoo Park, and Jeong Seop Sim. A fast algorithm for order-preserving pattern matching. *Inf. Process. Lett.*, 115(2):397–402, 2015.

[7] Maxime Crochemore, Costas S. Iliopoulos, Tomasz Kociumaka, Marcin Kubica, Alessio Langiu, Solon P. Pissis, Jakub Radoszewski, Wojciech Rytter, and Tomasz Walen. Order-preserving incomplete suffix trees and order-preserving indexes. In *SPIRE*, volume 8214 of *Lecture Notes in Computer Science*, pages 84–95. Springer, 2013.

[8] Peter Elias. Universal codeword sets and representations of the integers. *IEEE Trans. Information Theory*, 21(2):194–203, 1975. URL: `http://dx.doi.org/10.1109/TIT.1975.1055349`, `doi:10.1109/TIT.1975.1055349`.

[9] Simone Faro and M. Oguzhan Külekci. Efficient algorithms for the order preserving pattern matching problem. In *Algorithmic Aspects in Information and Management*, volume 9778, 2016. URL: `ttp://www.springerlink.com/index/10.1007/978-3-642-02158-9`, `doi:10.1007/978-3-642-02158-9`.

[10] Paolo Ferragina and Giovanni Manzini. An experimental study of a compressed index. *Inf. Sci.*, 135(1-2):13–28, 2001. URL: `http://dx.doi.org/10.1016/S0020-0255(01)00098-6`, `doi:10.1016/S0020-0255(01)00098-6`.

[11] Pawel Gawrychowski and Przemyslaw Uznanski. Order-preserving pattern matching with k mismatches. In *CPM*, volume 8486 of *Lecture Notes in Computer Science*, pages 130–139. Springer, 2014.

[12] Simon Gog, Timo Beller, Alistair Moffat, and Matthias Petri. From theory to practice: Plug and play with succinct data structures. In *SEA*, volume 8504 of *Lecture Notes in Computer Science*, pages 326–337. Springer, 2014.

[13] Simon Gog and Matthias Petri. Optimized succinct data structures for massive data. *Softw., Pract. Exper.*, 44(11):1287–1314, 2014. URL: `http://dx.doi.org/10.1002/spe.2198`, `doi:10.1002/spe.2198`.

[14] Roberto Grossi, Ankur Gupta, and Jeffrey Scott Vitter. High-order entropy-compressed text indexes. In *SODA*, pages 841–850. ACM/SIAM, 2003.

[15] D. Gusfield. *Algorithms on Strings, Trees, and Sequences: Computer Science and Computational Biology*. Cambridge University Press, 1997.

[16] Jinil Kim, Peter Eades, Rudolf Fleischer, Seok-Hee Hong, Costas S. Iliopoulos, Kunsoo Park, Simon J. Puglisi, and Takeshi Tokuyama. Order-preserving matching. *Theor. Comput. Sci.*, 525:68–79, 2014.

[17] Sebastian Kreft and Gonzalo Navarro. On compressing and indexing repetitive sequences. *Theor. Comput. Sci.*, 483:115–133, 2013.

[18] Marcin Kubica, Tomasz Kulczynski, Jakub Radoszewski, Wojciech Rytter, and Tomasz Walen. A linear time algorithm for consecutive permutation pattern matching. *Inf. Process. Lett.*, 113(12):430–433, 2013.

[19] Giovanni Manzini and Marcella Rastero. A simple and fast DNA compressor. *Softw., Pract. Exper.*, 34(14):1397–1411, 2004. URL: `http://dx.doi.org/10.1002/spe.619`, `doi:10.1002/spe.619`.

[20] Gonzalo Navarro and Veli Mäkinen. Compressed full-text indexes. *ACM Comput. Surv.*, 39(1), 2007. URL: `http://doi.acm.org/10.1145/1216370.1216372`, `doi:10.1145/1216370.1216372`.

[21] Gonzalo Navarro, Simon J. Puglisi, and Jouni Sirén. Document retrieval on repetitive collections. In *ESA*, volume 8737 of *Lecture Notes in Computer Science*, pages 725–736. Springer, 2014.

Appendix: Additional tables

$\|P\| = 15$		ibm	prices	tmax	ecg	rwalk	rand
ave # occs		92.63	59.70	1392.66	2142.49	2.38	1.00
$q = 4$	Phase 1	2.10	2.43	1.15	2.28	10.20	3.58
	Phase 2	1.01	1.04	1.00	1.07	4.04	2.02
$q = 8$	Phase 1	9.63	6.58	1.50	8.99	142.49	1.43
	Phase 2	1.02	1.03	1.00	1.04	2.63	1.00
$q = 12$	Phase 1	120.22	304.25	2.19	60.70	19105.57	436.38
	Phase 2	1.03	1.09	1.00	1.06	7.68	1.02

$\|P\| = 20$		ibm	prices	tmax	ecg	rwalk	rand
ave # occs		24.22	1.49	988.39	106.78	1.00	1.00
$q = 4$	Phase 1	2.42	1.13	1.05	3.31	1.14	1.00
	Phase 2	1.01	1.02	1.00	1.24	1.07	1.00
$q = 8$	Phase 1	7.87	42.19	1.44	21.75	1.26	1.00
	Phase 2	1.02	1.01	1.00	1.09	1.00	1.00
$q = 12$	Phase 1	31.38	140.24	1.59	93.16	39.29	1.00
	Phase 2	1.03	1.01	1.00	1.15	1.01	1.00

$\|P\| = 50$		ibm	prices	tmax	ecg	rwalk	rand
ave # occs		1.00	1.00	371.99	1.00	1.00	1.00
$q = 4$	Phase 1	1.00	1.00	1.02	1.00	1.00	1.00
	Phase 2	1.00	1.00	1.00	1.00	1.00	1.00
$q = 8$	Phase 1	1.00	1.00	1.07	1.00	1.00	1.00
	Phase 2	1.00	1.00	1.00	1.00	1.00	1.00
$q = 12$	Phase 1	2.18	1.00	1.12	1.00	1.00	1.00
	Phase 2	1.00	1.00	1.00	1.00	1.00	1.00

Table 7: False positives as a function of the window size q for the Strict Stock Market Index (ssmi). The first row shows the average number of actual occurrences for the patterns in the test set. The other rows show the ratios between candidates and actual occurrences at the end of Phase 1 and 2.